The Use of OpenSource in the Grid Computing and Data Acquisition System in High Energy and Particle Physics Research Projects

12th of October 2009 @ Paris - France

Richard Randriatoamanana
Computer Science Research Engineer and ICT Head Unit
IN2P3 Laboratory of Physics Nuclear and High Energy
CNRS and University of Pierre and Marie Curie
4, Place Jussieu - Paris 5e - France
randria@in2p3.fr

Fabio Hernandez
Deputy Director of IN2P3 Computing Center
Domaine scientifique de La Doua
43, bd du 11 Nov.1918 - 69622 Villeurbanne - France
fabio@in2p3.fr

Abstract

High-energy physics and astroparticle physics research has been traditionally performed by international collaborations. Open source software is an essential ingredient of the computing infrastructure required to support such research experiments at a global scale and in a development framework of an open collaboration. Components ranging from operating systems, data transfer systems, data processing and data visualization to project management and human collaboration tools.
In this presentation we will show examples of the usage of open source software and toolkits in two different contexts: a worldwide computing grid for processing the data of the Large Hadron Collider, the most powerful particle accelerator ever built and the largely distributed data acquisition and processing system built for the Pierre Auger Observatory.

1 Institutes and labs

1.1 CNRS

The **C**entre **N**ational de la **R**echerche **S**cientifique (National Center for Scientific Research) is a government-funded research organization, under the administrative authority of France's Ministry of Research. It was founded in 1939 by governmental decree, CNRS has the following missions:

- To evaluate and carry out all research capable of advancing knowledge and bringing social, cultural, and economic benefits for society.
- To contribute to the application and promotion of research results.
- To develop scientific information.
- To support research training.
- To participate in the analysis of the national and international scientific climate and its potential for evolution in order to develop a national policy.

The CNRS has 33,600 employees of which 26,000 are tenured employees: 11,600 researchers and 14,400 engineers and support staff. As the largest fundamental research organization in Europe, CNRS carried out research in all fields of knowledge, through its seven institutes:

- Institute of Chemistry (INC[1]) : develops and coordinates research into developing new compounds, understanding chemical reactivity, and elucidating and predicting relationships between the structure of compounds at the atomic level and their properties.
- Institute of Ecology and Environment (INEE[2]): develops and coordinates research carried out in the fields of global ecology : ecology and environment, including biodiversity and interactions between humans and environment.
- Institute of Physics (INP[3]): develops and coordinates research in the field of physics, with two main motives: to understand the world and to respond to current societal challenges.
- Institute of Biological Sciences (INSB[4]): develops and coordinates research in biological sciences with the overall aim to better understand life in all its complexity, starting at the level of atoms and biomolecules, extending to cells and through to organisms and populations.

[1]http://www.cnrs.fr/inc
[2]http://www.cnrs.fr/inee
[3]http://www.cnrs.fr/inp
[4]http://www.cnrs.fr/insb

- Institute for Humanities and Social Sciences (INSHS[5]): develops research on humans, from the development of language and knowledge to economic, social and political activities.
- Institute for Mathematical Sciences (INSMI[6]): develops and coordinates research in different branches of mathematics.
- Institute of Information and Engineering Sciences and Technologies (INST2I [7]): develops and coordinates research in the fields of information and engineering science and technology.

and two national institutes:

- National Institute for Earth Sciences and Astronomy (INSU[8]): develops and coordinates research with a national and international dimension in astronomy, earth sciences, ocean sciences, atmospheric and space sciences.
- National Institute of Nuclear and Particle Physics (IN2P3[9]): cf. *chapitre IN2P3*

CNRS encourages collaboration between specialists from different disciplines in particular with the university thus opening up new fields of enquiry to meet social and economic needs. CNRS has developed interdisciplinary programs which bring together various CNRS departments as well as other research institutions and industry. Interdisciplinary research is undertaken in the following domains:

- Life and its social implications
- Information, communication and knowledge
- Environment, energy and sustainable development
- Nanosciences, nanotechnologies, materials
- Astroparticles: from particles to the Universe

CNRS laboratories (or research units) are located throughout France, and employ a large body of tenured researchers, engineers, and support staff. Laboratories are all on renewable four-year contracts, with bi-annual evaluation by the National Center for Scientific Research. There are two types of labs:

- CNRS intramural labs: fully funded and managed by CNRS (called UPR, or *units propres de recherche*, in French)
- Joint labs: partnered with universities, other research organizations, or industry (called UMR, or *units mixtes de recherche*, in French)

[5]http://www.cnrs.fr/inshs
[6]http://www.cnrs.fr/insmi
[7]http://www.cnrs.fr/inst2i
[8]http://www.insu.fr
[9]http://www.in2p3.fr

The CNRS's annual budget represents a quarter of French public spending on civilian research. This funding comes from various sources:

- Government and public funding;
- CNRS funds, primarily from industrial and EU research contracts and royalties on patents, licenses, and services provided.

1.2 IN2P3

Created in 1971, the IN2P3's mission is to unite and promote research activities in the fields of nuclear physics, particle physics and astroparticle. It coordinates programs in these areas on behalf of CNRS and universities, in partnership with CEA[10]. This research aims to explore the physics of elementary particles and their fundamental interactions and their assemblies in atomic nuclei, to study the properties of these nuclei and to explore the connections between the infinitely small and infinitely large. If these themes are the heart of the discipline, IN2P3 also aims to contribute its expertise to other scientific fields, and the resolution of certain problems in society and participate in the training of students alongside the university. IN2P3 has finally the duty to benefit the business world for its expertise in the spreading technological resources to industrials issued by its research activities.
Scientific fields: particle physics, nuclear and hadron physics, astroparticles and neutrinos, research downstream of the electronuclear cycle, research and development for accelerators, computing grids...

1.3 IN2P3 Computing Center

It provides to its users, the tools to manage, store and process their scientific data. It provides also a whole of computing services destinated to the users, IN2P3 (CNRS) and IRFU (CEA) laboratory staff. The IN2P3 network, build over RENATER (french academic network) and the IN2P3 domain name (dns) are also managed by the Computing Center.

1.4 LPNHE

It's a *Unit Mixte de Recherche* of the National Institute of Nuclear Physics and Particle Physics (IN2P3), of CNRS and Universities Paris 6 and Paris 7. It employs 12 research groups, 3 technical services (computer science, electronics, mechanical), and 2 support services (administration, logistics). The laboratory is engaged in several large experimental research programs pursued in the context of international collaborations with very large research facilities around the world, centers of particle accelerators

[10]CEA is a French government-funded technological research organisation

and observatories. These programs cover current issues of particle physics, astroparticle and cosmology. The IT department ensures the development and administration lab's information systems, network, servers, data storage, messaging and user-support assistance. The unit is also involved in developments by coordinating experiments in acquisition systems, software and databases management. The laboratory houses a node of computing grid in the Ile-de-France, part of global grids that are LCG (LHC Computing Grid) and EGEE (Enabling Grids for E-sciencE): grid-based infrastructure, composed of hundreds of computing sites, linked by high-speed networks.

2 High Energy and Particle (HEP) Physic Experiments

2.1 CERN

We just cannot talk here about HEP Experiments without a short introduction about CERN[11], the European Organization for Nuclear Research, is one of the world's largest and most respected centres for scientific research. Its business is fundamental physics, finding out what the Universe is made of and how it works. At CERN, the world's largest and most complex scientific instruments are used to study the basic constituents of matter ? the fundamental particles. By studying what happens when these particles collide, physicists learn about the laws of Nature.
The instruments used at CERN are particle accelerators and detectors. Accelerators boost beams of particles to high energies before they are made to collide with each other or with stationary targets. Detectors observe and record the results of these collisions.
Founded in 1954, the CERN Laboratory sits astride the Franco-Swiss border near Geneva. It was one of Europe?s first joint ventures and now has 20 Member States.

2.2 LCG, the Large Hardon Collider Computing Grid

The Large Hadron Collider (LHC) is the largest scientific instrument on the planet at CERN near Geneva (Figure 1), where it spans the border between Switzerland and France about 100 m underground. It is a particle accelerator used by physicists to study the smallest known particles ? the fundamental building blocks of all things. It will revolutionise our understanding, from the minuscule world deep within atoms to the vastness of the Universe. Two beams of subatomic particles called 'hadrons' ? either protons or lead ions ? will travel in opposite directions inside the circular accelerator, gaining energy with every lap. Physicists will use the LHC

[11] http://cern.ch

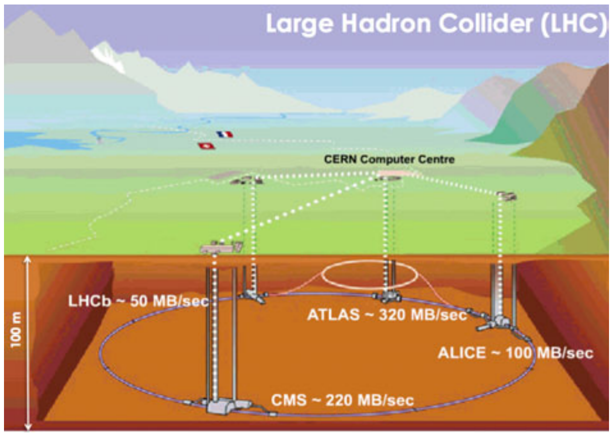

Figure 1: LHC detector topology

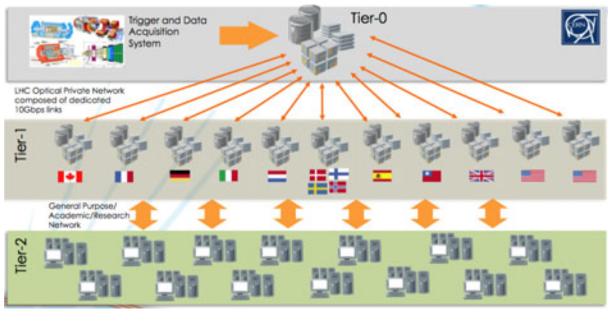

Figure 2: LCG Computing Connectivity

to recreate the conditions just after the Big Bang, by colliding the two beams head-on at very high energy. Teams of physicists from around the world will analyse the particles created in the collisions using special detectors in a number of experiments dedicated to the LHC.
When it begins operations, it will produce roughly 15 Petabytes (15 million Gigabytes) of data annually, which thousands of scientists around the world will access and analyse. The mission of the Worldwide LHC Computing Grid (WLCG) project is to build and maintain a data storage and analysis infrastructure for the entire high energy physics community that will use the LHC. The WLCG is a global collaboration of more than 170 computing centres in 34 countries. LCG-France is the french contribution to WLCG and this project is dedicated to the LHC computing launched in 2004 to provide a funded Tier-1 centre and an Analysis facility in the Computing Center of IN2P3 @ Lyon supporting the 4 LHC experiments (ALICE, CMS, LHCb, ATLAS) to promote the emergence of Tier-2 and even Tier-3 centres.
Computing resources arranged in a hierarchy of tiers with different roles and responsibilities

Figure 3: LCG Topology

2.3 Grid Computing in LHC

A Grid-based infrastructure (*figure 3*), composed of hundreds of computing sites, linked by high-speed networks. Integration of resources into a coherent environment that can be used by any collaboration member. A set of services and applications running on top of the grid infrastructures provided by EGEE[10] (Europe and Asia Pacific), Open Science Grid[11] (USA) and NorduGrid[12] (nordic countries). Operations of grid-based infrastructures:

- Supporting both end-users and site operators : GGUS
- Obtaining up-to-date information on the resources, sites, services and virtual organizations : operations portal site (cic.gridops.org)
- Accounting for resource utilization : Grid Operations Centre
- Sharing operation responsibility : operations portal site (cic.gridops.org)
- Coordinating handling of security incidents
- Certification of middleware and coordination of its deployment
- Monitoring of grid services :

Extremely Large Fabric management system (ELFms[9]) is the toolkit collection comprising - **QUATTOR**[9] for system management (configuration, installation, maintenance) - **LEMON**[9] for performance and exception monitoring - **LEAF**[9] HMS and SMS for CERN-CC Hardware and State Management - **SLS**[9] for Service Level Status.

2.4 AUGER/CDAS, the Pierre Auger Central Data Acquisition System

The history of cosmic ray research is a story of scientific adventure. For nearly a century, cosmic ray researchers have climbed mountains, ridden hot air balloons, and traveled to the far corners of the earth

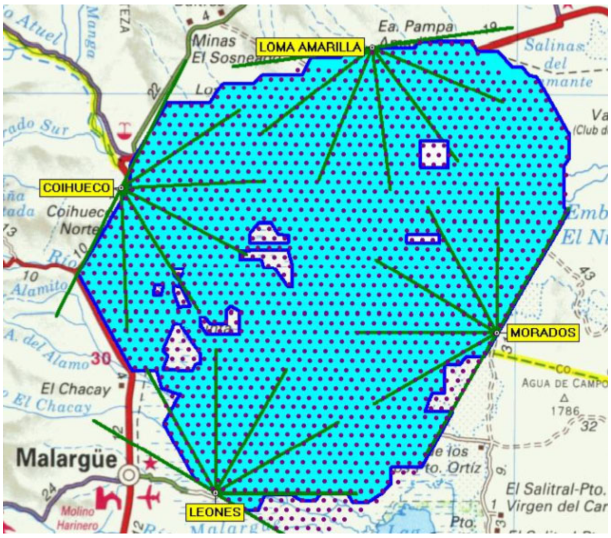

Figure 4: Pierre Auger Detector Array in Argentina

in the quest to understand these fast-moving particles from space. They have solved some scientific mysteries – and revealed many more. With each passing decade, scientists have discovered higher-energy, and increasingly more rare, cosmic rays. The Pierre Auger Project is the largest scientific enterprise ever conducted in the search for the unknown sources of the highest-energy cosmic rays ever observed.

The Observatory. On the vast plain known as the Pampa Amarilla (yellow prairie) in western Argentina, a new window on the universe is taking shape. There, the Pierre Auger Cosmic Ray Observatory is studying the universe's highest energy particles, which shower down on Earth in the form of cosmic rays. While cosmic rays with low to moderate energies are well understood, those with extremely high energies remain mysterious. By detecting and studying these rare particles, the Auger Observatory is tackling the enigmas of their origin and existence. The Auger Observatory is a "hybrid detector," employing two independent methods to detect and study high-energy cosmic rays. One technique detects high energy particles through their interaction with water placed in surface detector tanks. The other technique tracks the development of air showers by observing ultraviolet light emitted high in the Earth's atmosphere.
The first detection method uses the Observatory's main visible feature (Figure4) - the 1,600 water tanks that cover an enormous section of the Pampa and serve as particle detectors. Each 3,000-gallon (12,000 liter) tank, separated from each of its neighbors by 1.5 kilometers, is completely dark inside - except when particles from a cosmic ray air shower pass through it. These energetic particles are traveling faster than the speed of light in water when they

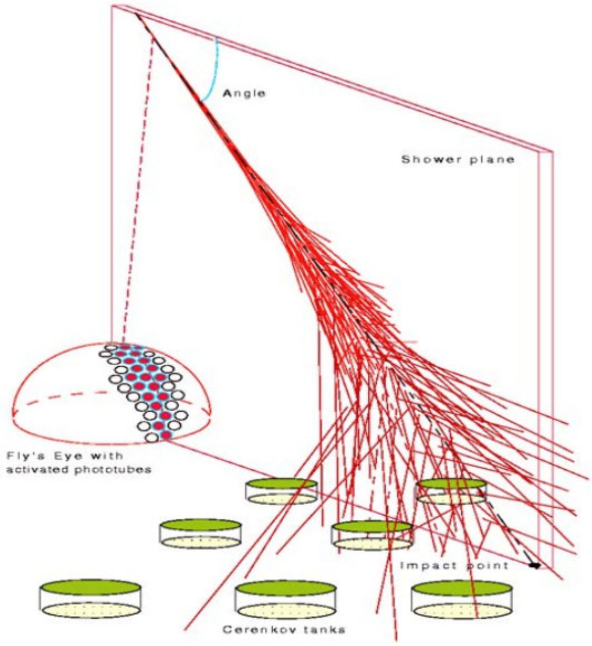

Figure 5: Pierre Auger Hybrid Detector

reach the detectors; therefore, their electromagnetic shock waves produce Cherenkov light that can be measured by photomultiplier tubes mounted on the tanks. Extensive air showers contain billions of secondary particles and can cause nearly simultaneous bursts of light in more than five tanks. Scientists can determine the energy of the primary cosmic ray particle based on the amount of light they detect from a sample of secondary particles. Slight differences in the detection times at different tank positions help scientists determine the trajectory of the incoming cosmic ray.

The Auger Central Data Acquisition System (CDAS) is the ICT heart of the experiment and is a custom built set of applications and its purpose is to operate the Surface Detector with minimal need of human intervention, provide interface with the Fluorescence Detector acquisition system for hybrid triggered events, and recollect, store, and prepare for transfer the auger data set.

2.5 Data acquisition system in AUGER, so called CDAS

The CDAS's daq/data flow processes (*figure 6*) is :

- an hybrid client/server architecture
- a postmaster process **to dispatch** raw/event data (**2Mb/sec** via TCP) to/from clients/array.
- a central trigger process **to decide** which tanks should be in a run array to send event raw data.

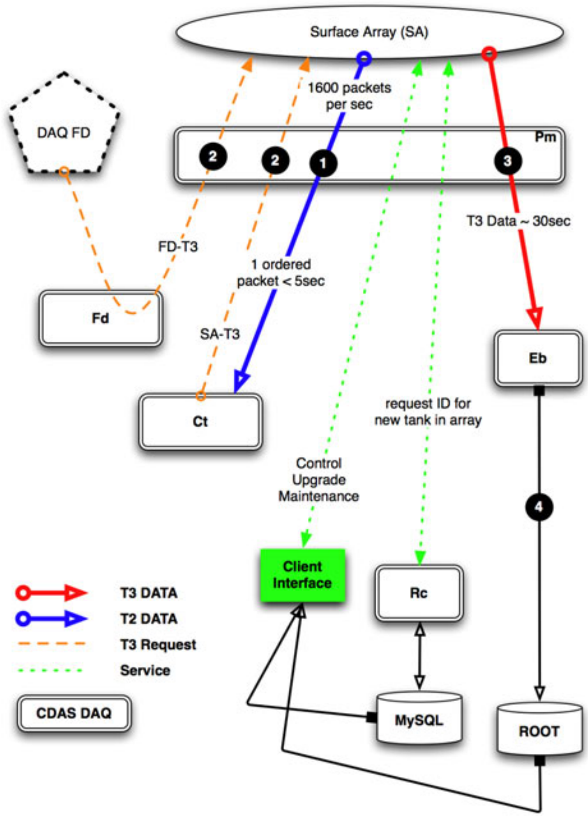

Figure 6: CDAS daq/data flow

- a event builder process **to build** from event raw data an auger event data stored in a database.
- a event displayer **to read** auger event data for analyze.

3 HEP common contribution in and uses of open-source

3.1 The web was born at CERN[3]

Tim Berners-Lee, a scientist at CERN, invented the World Wide Web (WWW) in 1989. The Web was originally conceived and developed to meet the demand for automatic information sharing between scientists working in different universities and institutes all over the world.
CERN is not an isolated laboratory, but rather a focus for an extensive community that now includes about 60 countries and about 8000 scientists. Although these scientists typically spend some time on the CERN site, they usually work at universities and national laboratories in their home countries. Good contact is clearly essential.
The basic idea of the WWW was to merge the technologies of personal computers, computer networking and hypertext into a powerful and easy to use

global information system.
Although the Web's conception began as a tool to aid physicists answer tough questions about the Universe, today its usage applies to various aspects of the global community and affects our daily lives.
Today there are upwards of 80 million websites, with many more computers connected to the Internet, and hundreds of millions of users. If households nowadays want a computer, it is not to compute, but to go on the Web.

3.2 Scientific Linux (SL), a linux tailor-made for scientific community

Extract from Scientific Linux web site[5]: *SL is a Linux release put together by Fermilab, CERN, and various other labs and universities around the world. Its primary purpose is to reduce duplicated effort of the labs, and to have a common install base for the various experimenters. The base SL distribution is basically [Red Hat] Enterprise Linux, recompiled from source. Our main goal for the base distribution is to have everything compatible with Enterprise, with only a few minor additions or changes. An example of of items that were added are Pine, and OpenAFS. Our secondary goal is to allow easy customization for a site, without disturbing the Scientific Linux base. The various labs are able to add their own modifications to their own site areas. By the magic of scripts, and the anaconda installer, each site is to be able to create their own distributions with minimal effort. Or, if a users wishes, they can simply install the base SL release.*

SLC is an SL variant that is built on top of the genuine SL and it is tailored to integrate within the CERN computing environment. SLC is fully compatible with SL therefore with the [Red Hat] Enterprise Linux: all software used or built on one of the versions should function properly on any other version.

..in numbers As of March 2009 around 36000 systems run Scientific Linux (SL), around 14000 systems run Scientific Linux CERN (SLC).

3.3 AFS[4], an open and remote FS widely deployed

AFS is an acronym for the Andrew File System, developed at Carnegie-Mellon University, Pittsburgh, under a sponsorship from IBM. For years, AFS has been maintained and sold by the Transarc Corporation, now part of IBM. Today, AFS has been "opened" by IBM, and is mainly maintained and distributed by the OpenAFS consortium.

AFS is a network-distributed file system comparable to Sun's NFS. AFS distinguishes between client machines and server machines. An AFS client enables users to access data residing on AFS servers transparently as if they were stored on a local disk. AFS servers in turn provide disk storage for files and directories.

AFS at CERN has been endorsed as the main networked file system for interactive Unix work. In certain cases however AFS is not suited and should not be relied on heavily. These cases include

- mission critical applications which must not be affected in case of network problems. AFS as a network file system relies (although not entirely) on a stable network;
- environments that include many systems for which AFS is not available, e.g. Lynx OS, OS/9. While it is possible to mount /afs on a non-AFS machine using NFS and a suitable AFS-gateway and therefore to improve interoperability, such a setup is more difficult to use and (as experience shows) not suited for 'heavy' traffic.

Stanford Linear Accelerator Center[6] uses AFS since 1998 and it supports all of the following projects in that most of the tools (compilers, CVS code trees, etc.) are in AFS as well as some of the data input and output for their simulations (some Monte Carlo type simulations but mostly data for GEANT simulations).

3.4 ROOT, a powerful data analysis framework

Definition extracted from ROOT website[8]: *The ROOT system provides a set of OO frameworks with all the functionality needed to handle and analyze large amounts of data in a very efficient way. Having the data defined as a set of objects, specialized storage methods are used to get direct access to the separate attributes of the selected objects, without having to touch the bulk of the data. Included are histograming methods in an arbitrary number of dimensions, curve fitting, function evaluation, minimization, graphics and visualization classes to allow the easy setup of an analysis system that can query and process the data interactively or in batch mode, as well as a general parallel processing framework, PROOF, that can considerably speed up an analysis.*

why ROOT? Having had many years of experience in developing the interactive data analysis systems PAW and PIAF and the simulation package GEANT, we realized that the growth and maintainability of these products, written in FORTRAN and

using some 20 year old libraries, had reached its limits. Although still very popular, these systems do not scale up to the challenges offered by the LHC, where the amount of data to be simulated and analyzed is a few orders of magnitude larger than anything seen before.
It became time to re-think our approach to large scale data analysis and simulation and at the same time we had to profit from the progress made in computer science over the past 15 to 20 years. Especially in the area of Object-Oriented design and development. Thus was born ROOT.
We started the ROOT project in the context of the NA49 experiment at CERN. NA49 generates an impressive amount of data, about 10 Terabytes of raw data per run. This data rate is of the same order of magnitude as the rates expected to be recorded by the LHC experiments. Therefore, NA49 was the ideal environment to develop and test the next generation data analysis tools and to study the problems related to the organization and analysis of such large amounts of data.
With ROOT we try to provide a basic framework that offers a common set of features and tools for all domains of High Energy Physics computing.
Currently the emphasis of ROOT is on the data analysis domain but thanks to the approach of loosely coupled object-oriented frameworks the system can easily be extended to other domains, like simulation, reconstruction, event displays and DAQ.

Architectural Overview. The backbone of the ROOT architecture is a layered class hierarchy with, currently, around 1200 classes grouped in about 60 frameworks (libraries) divided in 19 main categories (modules). This hierarchy is organized in a mostly single-rooted class library, that is, most of the classes inherit from a common base class TObject. While this organization is not universally popular in C++, it has proven to be well suited for our needs (and indeed for almost all successful class libraries: Java, Smalltalk, MFC, etc.). It enables the implementation of some essential infrastructure inherited by all descendants of TObject. However, we also can have classes not inheriting from TObject when appropriate (e.g., classes that are used as built-in types, like TString).

4 A (non-extensive) list of open-source tools used in HEP Computing Grid

Storage Resource Management[13] (SRM) It's a protocol for Grid access to mass storage systems (tape or disk or disk arrays). The protocol itself is a collaboration between Lawrence Berkeley (LBNL), Fermilab (FNAL), Jefferson (JLAB), CERN, and RAL

dCache provides a system for storing and retrieving huge amounts of data, distributed among a large number of heterogenous server nodes, under a single virtual filesystem tree with a variety of standard access methods

DPM (Disk Pool Manager) is a lightweight solution for disk storage management. If offers the required SRM interfaces, hopefully without being complicated by other modes of access or complications such as tape storage systems. It has been developed at CERN.

Grid Computing Middleware Softwares

IRODS stands for **i R**ule **O**riented **D**ata **S**ystems, is a project for building the next generation data management cyberinfrastructure.

SRB (Storage Resource Broker) provides a uniform interface to heterogeneous data storage resources over a network.

gLite distribution is an integrated set of components designed to enable resource sharing. In other words, this is middleware for building a grid.

References

[1] The Large Hadron Collider (LHC) at CERN

- http://lcg.in2p3.fr : LCG-France Official Website
- http://lcg.web.cern.ch/lcg/public : WLCG Official Website
- http://public.web.cern.ch/public/en/LHC/LHC-en.html : LHC Official Website

[2] Pierre Auger Observatory.

- `http://www.auger.org` : Official Website
- `http://www.auger.org.ar` : Southern Site
- `http://auger.in2p3.fr` : IN2P3 Auger's Official Website

[3] CERN, where the web was born.

- http://public.web.cern.ch/public/en/About/Web-en.html
- `http://info.cern.ch/` : The website of the world's first-ever web server

[4] AFS, a filesystem mostly implemented in scientific community.

- http://www.openafs.org
- http://consult.cern.ch/service/afs/afsguide.html
- http://www.ibm.com/developerworks/linux/library/os-afs.html

[5] Scientific Linux is a recompiled Red Hat Enterprise Linux, co-developed by Fermi National Accelerator Laboratory and the European Organization for Nuclear Research (CERN).

- `http://www.scientificlinux.org` : scientific linux portal
- `http://linux.web.cern.ch` : scientific linux for cern
- http://distrowatch.com/table.php?distribution=scientific

[6] SLAC National Accelerator Laboratory is home to a two-mile linear accelerator?the longest in the world.

- `http://www.slac.stanford.edu`

[7] Fermi National Accelerator Laboratory advances the understanding of the fundamental nature of matter and energy by providing leadership and resources for qualified researchers to conduct basic research at the frontiers of high energy physics and related disciplines.

- `http://www.fnal.gov`

[8] ROOT Cern.

- `http://root.cern.ch`

[9] CERN's Extremely Large Fabric management system : Lemon, Quattor, SLS and LEAF.

- `http://elfms.web.cern.ch/elfms` : Project ELFms
- `http://leaf.web.cern.ch` : LEAF
- `http://lemon.web.cern.ch` : Lemon
- `http://www.quattor.org` : Quattor

[10] http://www.eu-egee.org

[11] http://www.opensciencegrid.org

[12] http://www.nordugrid.org

[13] http://sdm.lbl.gov/srm-wg